JOINT DARK ENERGY MISSION (JDEM) PROJECT

Demonstration Science Data Processing and Quality Monitoring System

Stakeholder Requirements

June 30, 2010

Fermi National Accelerator Laboratory

Batavia, Illinois

REVISION STATUS

Version	Date	Changed By	Description
Original	6/30/10	Erik Gottschalk	Baseline

1 Introduction

1.1 Purpose of the Document

The purpose of this document is to specify stakeholder requirements for a demonstration science data processing and quality monitoring system for the Joint Dark Energy Mission (JDEM). JDEM is a collaboration between NASA and the US Department of Energy (DOE) to build a space telescope for the study of dark energy. The demonstration system will be developed at Fermilab for the JDEM Science Operations Center (SOC). The SOC is part of the JDEM Ground Data System. The demonstration system is referred to as the JDEM Demonstration Data Processing System (JDDPS).

1.2 Scope of the Software

The purpose of JDDPS is to provide science data processing capabilities for mission concept studies, demonstrate Fermilab capabilities for data processing, determine performance metrics, and provide a testbed to evaluate technologies for trade studies. The primary goal is to demonstrate a system that performs simulations and data processing for slitless spectroscopy and near-infrared (NIR) imaging for JDEM. The motivation is to provide better understanding of one of the JDEM science objectives, the study of Baryon Acoustic Oscillations.

As a science data processing and quality monitoring system, JDDPS will operate in a distributed computing environment. The system will consist of a prototype workflow with workflow participants that operate in this environment and encapsulate applications needed for slitless spectroscopy and NIR image processing. Substitute applications will be used in the workflow when actual applications are not available, so that an entire workflow can be assembled. One of the objectives is to study the performance of the workflow, which means that substitute applications need to be implemented so that they approximate the use of computational resources expected for actual JDEM applications. As new applications are developed for the workflow they will be incorporated into JDDPS.

The scope of JDDPS includes the development of simulation, data processing, and quality monitoring software for JDEM slitless spectroscopy and NIR image processing. Simulation and data processing will be implemented as a prototype workflow that includes rudimentary workflow management and provenance tracking capabilities. Workflow capabilities will be implemented by using the Kepler workflow system. This amounts to an evaluation of Kepler, which is therefore included in the scope of JDDPS. Also included in the scope is the evaluation of OpenSplice DDS message passing software for quality monitoring, and the use of databases for slitless spectroscopy science data processing. Each of these evaluations will be performed within the context of JDDPS, and suitability of the software will take into account the level of effort needed to satisfy the requirements presented in this document.

1.3 Stakeholders

This stakeholder requirements document for JDDPS is intended for the following stakeholders:

- DOE's JDEM Project Office,

- DOE's JDEM GDS team,
- DOE's JDEM Scientists,
- Fermilab Computer Security Team, and
- Fermilab Management.

It is worth noting that the JDEM Interim Science Working Group (ISWG) is not listed as a stakeholder, but members of the ISWG are included as stakeholders through the JDEM Project Office and as JDEM scientists and are therefore able to influence the development of JDDPS.

DOE's JDEM GDS team and JDEM scientists include both software developers and users of JDDPS. The developers and users are organized according to roles that identify different types of functional requirements. We define four roles:

- Scientist algorithm developer - A scientist with specific expertise in developing and implementing algorithms for processing JDEM data.
- Scientist pipeline developer - A scientist with expertise in developing and implementing data processing pipelines to produce data products.
- Scientist data analyst - A scientist who uses data products to perform scientific analyses of data.
- Operator - A person who operates data processing pipelines.

Three of the roles are usually (but not exclusively) associated with work done by scientists, so we use the term "scientist" to define these roles. The "operator" role is often performed by a scientist, but we do not view this role as requiring much of a background in science. Other roles associated with software development, such as developing software infrastructure, are not included in this stakeholder requirements document, but are important for system requirements.

1.4 Definitions, Acronyms, and Abbreviations

The following definitions, acronyms, and abbreviations provide the meanings for terms used in this document.

Actor - A software component that performs a task, typically by reading input and producing output.

API - Application Programming Interface

BAO - Baryon Acoustic Oscillations

Campaign - A workflow initiated by a human.

DDS - Data Distribution Service, a customizable quality of service publish/subscribe standard from the Object Management Group.

DOE - Department of Energy

GDS - Ground Data System

JDEM - Joint Dark Energy Mission

Instantiated workflow - A workflow in which all data sources and participants, as well as their configurations, are specified.

Job - A submission to a batch processing queue.

NASA - National Aeronautics and Space Administration

NIR - Near Infrared

OMG - Object Management Group, a group that maintains standards for distributed, object-oriented software systems.

Participant - An actor whose action is triggered by a workflow engine. From the point of view of the workflow engine, the task carried out by a participant is atomic in that the task completes successfully or fails to complete. Moreover, the workflow engine does not manipulate the internal state of a participant. A participant might, for example, be a shell script that runs a data processing application to perform a specific task.

Pipeline - A workflow in which the participants are arranged according to a pipe and filter architecture. In such an architecture, the participants process units of work and can execute concurrently, each reading input from its predecessor and providing output to its successor.

Provenance - A record that identifies the participant that was responsible for creating a data product.

Publisher - A software entity that prepares data for transmission based on one or more writers.

QC - Quality Control

QoS - Quality of Service, is the ability to provide different priorities to data flow in a network. QoS refers to control mechanisms that are used to reserve resources for different applications, users, or data flows, or to guarantee a certain level of performance.

Reader - A software entity that reconstitutes an object from the data that has been received from a writer.

SOC - Science Operations Center

Stream of data - A sequence of units of work of the same type.

Submit node - A node that manages a job and can monitor other nodes involved in executing a job.

Subscriber - A software entity that receives data using one or more readers.

Unit of work - The smallest data element that is processed in its entirety by an actor. Usually actors operate on a sequence of data elements.

Virtual machine - A virtual machine is a software implementation of a computer that executes programs like a physical computer.

Workflow - A collection of participants and a defined set of rules specifying when (under what conditions) and how (with what parameters, configuration, and input data) the actions performed by participants should be triggered.

Workflow engine - A software application that manages and executes workflows.

Workflow management system - A software application that triggers the actions performed by participants according to the rules that define the workflow. A workflow management system may additionally record provenance and other metadata about the execution of the workflow, and may provide tools for users to specify workflows.

Workflow template - A workflow in which a subset of the data and/or participants are specified abstractly. At a minimum, a workflow template contains only the workflow rules, which refer to data and actors using undefined symbols.

Writer - A software entity that makes instances of a user-defined data structure available over a network.

1.5 References

This requirements document shall be used in conjunction with the documents listed in this section. URLs are provided for documents that are available. Documents that are not yet available are under review and will be released as soon as they become available. When documents are superseded by an approved revision, the revision shall apply.

1) JDDPS Concept of Operations - currently under review

2) QuIDS Concept of Operations - http://cdcvs.fnal.gov/redmine/documents/show/61

1.6 Overview of the Document

This document specifies stakeholder requirements for JDDPS. Section 2 provides a general description of JDDPS. It describes the relationship of JDDPS to other projects, functions and purpose, operating environment, and general constraints. Section 3 stipulates the requirements, and is organized into several subsections. The subsection titled "Functional Requirements" is organized according to the four roles defined in Section 1.3 for different types of software developers and users. Each role corresponds to a set of functions that will be provided by JDDPS.

2 General Description

2.1 Relation to Current Projects

JDDPS benefits from its relationship to three projects at Fermilab and one project in which Fermilab is collaborating with Tech-X Corporation (http://www.txcorp.com/). The projects are LQCD, NOvA, QuIDS, and the FermiCloud project. The relationship of these projects to JDDPS is described in the following paragraphs.

LQCD is a project that uses the Lattice QCD (quantum chromodynamics) approach to solving theoretical problems in particle physics. LQCD faces significant computational challenges and has often been at the forefront of developing high-performance computing capabilities. At Fermilab the work has included a recent evaluation of workflow management systems (see JDDPS ConOps document referenced in Section 1.4), and has influenced our approach to workflow management for JDDPS.

The NOvA project is a particle physics experiment that is being constructed to detect and analyze neutrino particles and their interactions. Members of the NOvA team have worked with the JDEM GDS team to evaluate specific implementations of the Data Distribution Service (DDS) message passing system. The NOvA team concluded that OpenSplice DDS satisfies their requirements for the NOvA data acquisition system, and has proceeded to implement functions using OpenSplice DDS software. The JDEM GDS team has not completed its evaluation, but will do so as part of the development of quality monitoring capabilities for JDDPS. If DDS satisfies JDEM requirements, then JDEM may benefit by using software developed for NOvA. For example, a message logger that is being developed for NOvA may be applicable to JDEM.

Members of the GDS team have been involved in the evaluation of a message passing system for quality control (QC) data. The name of the system is "Quality Information Distribution Service" (QuIDS), and it is based on DDS. QuIDS is being developed together with Tech-X, which received a Phase-1 SBIR grant to evaluate different DDS implementations and to develop a test system to measure performance limits of the messaging software. Having successfully completed Phase 1 of the SBIR, Tech-X has submitted a proposal for a Phase-2 grant to implement QC capabilities for JDEM.

JDEM has contributed computing hardware to the FermiCloud project, which aims to investigate and implement "cloud computing" capabilities at Fermilab. The computing hardware will be used to run tests and demonstrations of JDDPS. The GDS team benefits from this collaboration since hardware installation and maintenance is managed by the FermiCloud project.

2.2 Relationship to Predecessor and Successor Projects

JDDPS benefits from expertise of JDEM GDS team members with previous experience from the Sloan Digital Sky Survey (SDSS), a survey that shares many of the features of JDEM. Members of Fermilab's Experimental Astrophysics Group (EAG) have extensive experience with data collection, processing, distribution and analysis of SDSS data. The EAG hosted and operated those portions of SDSS science data processing infrastructure that corresponds to infrastructure that needs to be developed for JDEM. Furthermore, much of the SDSS infrastructure was developed by the EAG.

One of the important contributions that the EAG made to SDSS was the workflow system. The SDSS workflow system was adequate for its time and satisfied the requirements for SDSS science data processing. However, the system is outdated and specific to SDSS infrastructure. Therefore, it is not suitable for JDDPS or for JDEM.

2.3 Functions and Purpose

The purpose of JDDPS is to demonstrate science data processing capabilities for slitless spectroscopy and NIR image processing for JDEM. JDDPS functions will be developed to satisfy the goals identified in the JDDPS ConOps document referenced in Section 1.4. The goals are reiterated here:

- Develop slitless spectroscopy simulations and data analysis capabilities for JDEM to support mission concept studies and assist in the development of data processing algorithms and software infrastructure.

- Develop a prototype workflow for slitless spectroscopy and NIR science data processing.

- Develop rudimentary workflow management and provenance tracking capabilities for the prototype workflow system.

- Evaluate the OpenSplice implementation of the Data Distribution Service standard for data processing quality control.

- Evaluate the use of databases for slitless spectroscopy and NIR science data processing.

2.4 Environmental Considerations

The Fermilab Computing Division supplies and/or manages computing systems for user access, application development, production processing, and analysis of experimental data. The production processing infrastructure consists of large clusters (many thousands of computing cores) of either loosely or tightly interconnected worker nodes, shared among various experimental groups but allocated according to needs and dedicated investments. The FermiCloud infrastructure is currently being deployed to provide resources for virtualized application services in support of experimental data processing. JDEM has contributed computing hardware to the FermiCloud project, and JDDPS is being designed to operate on FermiCloud nodes.

The Computing Division also supplies and/or operates data storage facilities and databases for experimental data. JDEM has purchased two database servers for JDDPS, and these will be managed by members of the Computing Division.

2.5 General Constraints

JDDPS does not involve any aspects beyond the usual and customary issues of running data processing systems at Fermilab. JDDPS will comply with Fermilab's Computer Security Policy. There are general constraints in that JDDPS must use existing hardware and services that exist at Fermilab. There are resource constraints in that users of JDDPS will share resources with other users at Fermilab when running on FermiCloud nodes.

3 Specific Requirements

Requirements for JDDPS are presented in this section. We use the term "system" to refer to the JDDPS demonstration system. The term "production system" refers to an instance of JDDPS running at Fermilab with all of its functions implemented. We use the term "development environment" to refer to the software development tools that are used to develop, build and test software for JDDPS. We identify two development environments. The "pipeline development environment" is used by "scientist algorithm developers" and "scientist pipeline developers" (see definition of roles in Section 1.3) to develop and implement algorithms and pipelines for JDEM. For a scientist or software developer who needs to use this development environment at their home institution they must satisfy requirements for the development environment. The second development environment is the "system development environment," which is used by system engineers to develop software infrastructure for JDDPS.

3.1 Functional Requirements

3.1.1 Scientist Algorithm Developer

Requirement STK.Req : 1
The pipeline development environment shall have a test suite to validate participants.

Requirement STK.Req : 2
Participants that satisfy the test suite shall be executable in a production system.

3.1.2 Scientist Pipeline Developer

Requirement STK.Req : 3
The pipeline development environment shall have a test suite to validate pipelines.

Requirement STK.Req : 4
Pipelines that satisfy the test suite shall be executable in a production system.

Requirement STK.Req : 5
The system shall have a method for establishing dependencies between participants.

Requirement STK.Req : 6
The system shall have a method for specifying input and output data products for participants.

Requirement STK.Req : 7
The system shall have a method for specifying the configuration of each participant.

3.1.3 Scientist Data Analyst

Requirement STK.Req : 8
A production system shall provide provenance for each data product that it creates.

Requirement STK.Req : 9
For each data product, the system shall provide a method for identifying all derived data products.

3.1.4 Operator

Requirement STK.Req : 10
The system shall execute user-submitted campaigns in a production system.

Requirement STK.Req : 11
The system shall provide the ability to run multiple campaigns simultaneously.

Requirement STK.Req : 12
The system shall provide a method for configuring campaigns.

Requirement STK.Req : 13
The system shall provide current status information for campaigns.

Requirement STK.Req : 14
The system shall provide a list of all active campaigns for a given user.

Requirement STK.Req : 15
The system shall provide access to quality control products for each active campaign.

Requirement STK.Req : 16
The system shall provide access to scientific data products for each active campaign.

Requirement STK.Req : 17
The system shall provide the means for immediate termination of an active campaign.

Requirement STK.Req : 18
The system shall provide the means for orderly premature termination of an active campaign.

Requirement STK.Req : 19
The system shall provide the means to restart a prematurely terminated campaign beginning with the last participant that completed a unit of work.

Requirement STK.Req : 20
The system shall retain data products created by a production system.

Requirement STK.Req : 21
The system shall provide remote active display of user-selected quality control products.

Requirement STK.Req : 22
The system shall provide a list of quality control products that are available for display.

Requirement STK.Req : 23
The system shall provide a method to change the selection of quality control products for display during the operation of a production system.

Requirement STK.Req : 24
The system shall provide a method to change the configuration of how quality control products are displayed during the operation of a production system.

Requirement STK.Req : 25
The system shall provide access to quality control products for all previous campaigns.

Requirement STK.Req : 26
The system shall provide a method for users to request notifications of job failures.

3.2 Performance Requirements

Requirement STK.Req : 36
The system shall process the data products that constitute a 24-hour period of JDEM data in less than 24 hours.

3.3 Interface Requirements

Requirement STK.Req : 37
The system shall be able to read FITS files.

Requirement STK.Req : 38
The system shall be able to write data products as FITS files.

3.4 Operational Requirements

Requirement STK.Req : 27
The system shall provide repositories for all software artifacts.

Requirement STK.Req : 28
The system shall provide access to repositories to all authorized collaborators.

Requirement STK.Req : 29
The system shall use validated releases for running demonstrations and data challenges.

Requirement STK.Req : 30
The system shall comply with policies established for Fermilab's operational environment.

3.5 Resource Requirements

JDDPS will operate in the FermiCloud environment. FermiCloud is a shared resource that JDEM has contributed to, and therefore JDEM users can expect to gain access to computing resources. Users of FermiCloud access computing resources by requesting access to virtual machines. For JDEM we define three types of requests (small, medium and large) to run JDDPS campaigns. Small, medium, and large campaigns are defined as requiring a maximum of 8, 16, and 32 virtual machines, respectively.

Requirement STK.Req : 31
JDEM users shall be able to run multiple small campaigns per week with aggregate execution time per node not exceeding 12 hours.

Requirement STK.Req : 32
JDEM users shall be able to run multiple medium campaigns per week with aggregate execution time per node not exceeding 24 hours.

Requirement STK.Req : 33
JDEM users shall be able to run multiple large campaigns per month with aggregate execution time per node not exceeding 24 hours.

3.6 Verification Requirements

JDDPS is being developed using best practices encouraged by the Project Management Institute (http://www.pmi.org/). Requirements are being developed and managed using IBM Rational DOORS. DOORS is also being used for validation.

Requirement STK.Req : 39
The system shall have a test suite to validate requirements.

3.7 Acceptance Testing Requirements

Requirement STK.Req : 40
The system shall have a validation procedure for stakeholder acceptance.

3.8 Security Requirements

Requirement STK.Req : 34
The system shall comply with policies established for Fermilab computing security.

3.9 Safety Requirements

Requirement STK.Req : 35
The system shall comply with Fermilab safety policies.